

The Chomsky–Place Correspondence 1993–1994

Noam Chomsky

Massachusetts Institute of Technology

Ullin T. Place

University of Wales, Bangor

Edited, with an Introduction and
Suggested Readings, by Ted Schoneberger
California State University, Stanislaus

Edited correspondence between Ullin T. Place and Noam Chomsky, which occurred in 1993–1994, is presented. The principal topics are (a) deep versus surface structure; (b) computer modeling of the brain; (c) the evolutionary origins of language; (d) behaviorism; and (e) a dispositional account of language. This correspondence includes Chomsky's denial that he ever characterized deep structure as innate; Chomsky's critique of computer modeling (both traditional and connectionist) of the brain; Place's critique of Chomsky's alleged failure to provide an adequate account of the evolutionary origins of language, and Chomsky's response that such accounts are "pop-Darwinian fairy tales"; and Place's arguments for, and Chomsky's against, the relevance of behaviorism to linguistic theory, especially the relevance of a behavioral approach to language that is buttressed by a dispositional account of sentence construction.

In 1996 I presented a paper (Schoneberger, 1996) on Noam Chomsky at the Association for Behavior Analysis annual convention in San Francisco. Ullin T. Place was in the audience. Afterwards, Ullin and I had a brief conversation about Chomsky; during that conversation I learned that he had corresponded extensively with Chomsky. At the end of our conversation Ullin promised to send me a copy of that correspondence. A few days after returning home from the convention I was delighted to receive, in the mail, a copy of the Chomsky–Place letters on a computer disk.

In January 2000 I learned of Ullin's passing from a posting on the Verbal Behavior Special Interest Group listserv hosted by Bill Potter. I responded by posting a message informing the listserv members that I had a copy of letters exchanged between Ullin and Noam Chomsky. Several members expressed an interest in seeing the cor-

respondence, including, especially, Hank Schlinger, editor of *The Analysis of Verbal Behavior (TAVB)*. The possibility of publishing them in *TAVB* was raised. Through the good offices of David Palmer, the Place family agreed to the publication of Ullin's letters. However, Chomsky indicated a reluctance to see his letters published. He reported that he engages in a large, time-consuming correspondence. As a consequence, his letters are often written quickly and in a casual manner. In addition, Chomsky indicated that his letters often contain comments of a personal nature that are inappropriate for publication.

In the ensuing months Hank Schlinger negotiated with Chomsky in an attempt to gain permission to publish his (Chomsky's) letters. In the end, Chomsky graciously agreed to allow an edited version of his letters to be published. Specifically, a small number of personal references were deleted by Chomsky. In addition, some typographical (and other) minor errors were corrected by him. Chomsky's deletions

Address correspondence to Ted Schoneberger, P.O. Box 157, Turlock, California 95381.

and corrections are not recorded in the text by means of standard editing devices (e.g., ellipsis points, bracketing). Similarly, I have deleted some paragraphs, and have made some additional minor corrections, without using standard editing protocol to record those changes. However, all other modifications I have made (e.g., adding substantive words, references to published works) do employ standard editing devices. I feel confident that none of these modifications alter the meaning of the letters.

Excerpts from four letters by Place are included. These are dated June 24, 1993; August 21, 1993; December 29, 1993; and May 16, 1994. Four letters by Chomsky are also included; they are dated February 24, 1993; July 22, 1993; November 8, 1993; and January 18, 1994. There are five major topics discussed in the letters: (a) deep versus surface structure; (b) computers and the brain; (c) pop-Darwinian fairy tales; (d) behaviorism; and (e) dispositions. These letters have been edited so that exchanges between Chomsky and Place concerning each specific topic are grouped together. For example, Chomsky's comments, in his initial letter, on deep versus surface structure are followed by Place's comments on this subject in his responding letter, followed by Chomsky's comments in his next letter, and so on. Each topic is so treated in turn.

Prior to the exchanges, Place had sent a number of his papers to Chomsky for comment. The exchanges thus begin with Chomsky responding to an assertion made by Place in one of these papers (i.e., Place, 1992). Specifically, against Place, Chomsky denies that he ever characterized deep structure as *innate*. This constitutes Chomsky's opening salvo in the debate over deep versus surface structure. Other topics follow. In the discussion of computers and the brain, Chomsky argues against the utility of both traditional and connectionist computer modeling of the brain. The discussion of behaviorism centers on Chomsky's maintaining its

irrelevance to linguistic theory, while Place offers spirited and elaborate argument for its relevance—including, especially, a behavioral approach that includes a dispositional account of sentence construction.

DEEP VERSUS SURFACE STRUCTURE

Chomsky (February 24, 1993)

[In your paper "Eliminative Connectionism" (Place, 1992, p. 22)] you refer to "Chomsky's notion of novel sentences being generated by a set of in built syntactic rules whose 'deep structure' is innate." The same comment is repeated several times. . . . But I've never held such a doctrine. In the earliest work of mine in generative grammar, from the 1940s into the early 1960s, there is no notion of "deep structure" at all. The concept was introduced, as a technical idea of no major significance, in work of the mid-1960s (in one form, in my *Aspects of the Theory of Syntax*) [Chomsky, 1965]. I never suggested that "deep structure" was innate, or that other aspects of language were not. Evidently, many people were misled by the term "deep," assuming it means something profound; in fact, "deep structure" is no more profound than "surface structure." To try to overcome these misinterpretations, I began, in the 1970s, to replace "deep" and "surface structure" by "D-structure" and "S-structure," hoping that people would not be misled by what is so obviously a technical notion. I'm afraid that the confusion persists to the present, even though, in my most recent work, I've found some interesting ways to reconstruct tentative ideas of the 1950s, which it wasn't possible to realize at that time, and to dispense with "deep structure" entirely—while, at the same time, vastly increasing what is taken to be innate, as seems to be required by simple empirical considerations.

Place (June 24, 1993)

Thank you for putting me straight concerning the point in your intellec-

tual development when you introduced the distinction between “deep” and “surface” structure. However, I am somewhat mystified by your denial that knowledge of deep structure should be seen as innate, in contrast to knowledge of surface structure which, on any account, must presumably be taken as learned. You do not deny, I take it, that on your view knowledge of deep structure accounts for linguistic universals, while knowledge of surface structure accounts for those features which vary from natural language to natural language. Nor, it would seem, do you deny that, from the mid-1960s at least, you have argued in favor of a nativist/maturational theory of the acquisition of linguistic competence. If so, it is hardly surprising that commentators like myself should assume that the part of linguistic competence that is innate is knowledge of deep structure and the part that is learned is knowledge of surface structure. If you concede, as you now apparently do, that part of our knowledge of deep structure may also be learned, you open the door, so it seems to me, for someone like myself who wants to argue that *all* linguistic competence is learned and explain the (“deep”) structure that is common to all natural languages on the assumption that every natural language is in the business of depicting environmental reality by constructing sentences whose syntactic structure reflects the structure of the segment of reality that is thereby depicted.

Chomsky (July 22, 1993)

I have never, at any time, suggested that deep structure is innate and surface structure learned, that deep structure accounts for linguistic universals and surface structure not, etc. The terms were introduced in work of the mid-1960s (among them, *Aspects of the Theory of Syntax* [Chomsky, 1965]), none of which even hinted at these conclusions in any way. At no point was it suggested, or believed, that these distinctions correlate with in-

nate-learned in anything remotely like the way you describe. The same is true of the notions they replaced, T-marker and phrase-marker (in the framework of earlier work). In work of the past few years, I’ve been suggesting that deep structure doesn’t exist at all; it was an error, and we should return to a revised version of the theory of generalized transformations. Does that imply that I am making the (quite absurd) assumption that every aspect of language is “learned” (whatever “learning” may be)? Surely not.

These are common misunderstandings outside of the technical disciplines, misunderstandings so deep-rooted that I finally just abandoned the terms, realizing that it was hopeless to try to correct the misconceptions, which apparently derive from some connotations of “deep” and “surface.” That’s why, years ago, I began using the terms “D-structure” and “S-structure,” hoping that that would not lead to the same gross misreadings.

Take, say, what was later called “binding theory,” a major subject for the study of universals. It’s been pretty clear (at least to me) from the mid-1960s that these principles did not hold of deep structure; it’s generally been held that they hold of surface structure. Or, take the work that Halle and I were doing on phonology (together, from the mid-1950s, separately before), leading to *Sound Pattern of English* [Chomsky & Halle, 1968]. Our primary interest was universals, but they all involve the parts of derivations leading from surface structure to phonetic form.

You say I seem to “concede” that deep structure might be partly learned. The word is a curious choice: I’ve always *insisted* on that point, while also *insisting* that most of the rest of the computational system, down to phonetic and logical form, is not learned and could not be. Furthermore, these points have never even been controversial within the field; everyone assumes that this might well be the case, and most have also assumed that it *is* the case. It has absolutely nothing to

do with the nativist/maturational issue. Furthermore, there *is no* “nativist/maturational theory” of language acquisition, any more than there is a “nativist/maturational theory” of the growth of any other aspect of the organism, and for the same reason. There are particular theories, all but the maddest of them “nativist/maturational,” exactly as in the case of a chicken’s wings. That’s exactly why the controversy over the “innateness hypothesis” has had such a curious character. There are many criticisms of this hypothesis, ranging up to quite furious, and no defense of it, none at all. The reason is, simply, that there is no such hypothesis, any more than there is an “innateness hypothesis” concerning the visual system; rather, specific hypotheses, all of them assuming innateness, naturally.

Place (August 21, 1993)

I must confess to being totally confused about the nature of your distinction between deep and surface structure. You appear to be, not just conceding, but *insisting* that most, if not all, of our knowledge of deep structure is learned and at the same time (a) denying that deep structure accounts for linguistic universals and (b) “insisting that most of the rest of the computational system, down to phonetic and logical form, is not learned and could not be.” I find this last sentence particularly puzzling partly because your subsequent repudiation of the computer model makes it very difficult to understand what you mean by the phrase “computational system” and partly because, assuming that by “the computational system” you mean the process in the brain whereby a speaker/listener construes, constructs, and articulates intelligible sentences, it is far from clear what “most of the rest of the . . . system” which “is not learned and could not be” consists in. You cite “phonetic and logical form.” Does this mean that you are asking us to believe that the speaker/listener does not and could not learn to recognize and pro-

duce the different patterns of sound which make up those words he or she is able to recognize and produce? Are you denying that anyone does or could learn to respond to and make use of the distinction between saying *All men are mortal* and *Some men are mortal*. *Others are immortal*? If this *is* what you are saying, I fail to understand why you should suppose that such things are not and could not be learned. If it is *not* what you are saying, I simply don’t know what it is that you are saying. Not only do I fail to understand the distinction between S-structure and D-structure, I have no grasp whatever of this apparently crucial innate residuum, the rest of the “system” which constitutes the individual’s linguistic competence, but which does not apparently consist in knowledge of the S- and D-structures.

Chomsky (November 8, 1993)

I do not see what puzzles you. The statements you quote are clear and accurate. I suspect that you are starting with some assumptions about the meanings of all these terms, and their place in linguistic theory, that is not accurate. Let me try once again to clarify, from a different point of view.

Right now I happen to be lecturing on a conception of language that works something like this, cutting corners for brevity. The (relevant parts of the) language faculty have an initial state S_0 , determined by biological endowment; the theory of this state is UG [universal grammar]. The state can be described in various terms (atoms, cells, C-R [computational-representational] systems that articulate its properties, etc.). I’m keeping here to C-R systems, in part because that is the only account which, for the moment, has any particular status by the criteria of normal science. Considered in these terms, S_0 is a “principles and parameters” system in the sense of work of the past 15 years or so [see Chomsky, 1997]. That is, it incorporates a system of principles and of (probably two-valued) pa-

parameters with values to be set by early experience. The principles distinguish a lexicon from a computational system. The parameters are very likely limited to the lexicon, in fact, to a small subpart of the lexicon involving functional features. Apart from parameter-setting, language learning involves settling such peripheral matters as Saussurean arbitrariness (do we use “tree” or “Baum”) and the like.

The computational system, C, is invariant. In this sense, there is only one language, as a rational Martian would have expected. C maps an array A of lexical choices (I leave open its exact structure; take it to be a sequence, for concreteness, though interesting issues are concealed here) to symbolic objects at the interface of the language faculty and other systems of the brain. It appears that there are two such interfaces: articulatory-perceptual, conceptual-intentional (a lot more is known about the first). Therefore C maps A to a pair $\langle \text{PF}, \text{LF} \rangle$ [phonetic form, logical form], the “instructions” for performance. A person who has the language, and is otherwise unimpaired, can (sometimes) use the instructions to articulate, perceive, ask questions, refer to things, etc.

What is “innate” is the general architecture, the specific principles, the general computational conditions (e.g., of economy) that carry the huge empirical burden of providing an account and explanation for the vast variety of empirical phenomena, including typological distinctions (why are the properties of Icelandic or Japanese so different from those of English), etc. In particular, C and its manner of functioning seems to be entirely innate. What is not innate is the choice of parameter values; e.g., universally, the case system appears to be an invariant part of C, but languages differ as to whether computations involving case feed the articulatory-perceptual system (Latin, Finnish, Sanskrit in somewhat different ways) or not (Chinese, English virtually, etc.). Similarly, question-formation seems to be a universal

part of C, but languages differ as to whether placing of the question word in a peripheral (scopal) position feeds the articulatory system (as in English) or not (as in Chinese).

Note that I have said nothing here about deep structure and surface structure. The reason is that I believe that they do not exist. But I am saying a great deal about innateness, in fact, far more than was claimed in earlier theories that did postulate the existence of deep structure (as an internal interface between the lexicon and C) and surface structure (as a designated internal level, with specific properties and conditions, at a point where computation “splits” to PF and LF). Hence it simply cannot be that the question of innateness has anything special to do with deep structure, nor did I ever imagine otherwise.

Whether deep and surface structure exist is an empirical question; it seemed 20 years ago that they did; it seems now (to me at least) that they don’t. Either way, the basic structure of language must be innate, and we now know a lot (much more than 20 years ago) about what that innate structure is, and how remarkably restrictive it is (more so, in fact, under the assumptions of theories like the one sketched, which abandon deep and surface structure entirely).

Suppose we go back, say, 20 years. I then assumed that these levels did exist, as a matter of fact. I also assumed that some of the properties of deep structure were innate, some learned; and some of the properties of surface structure were innate, some learned. Thus I assumed that languages could vary as to whether they had functional categories at deep structure (in fact, I still assumed that a few years ago, and the analogue of that belief in the newer system also remains an open question). Or whether they had ordering differences at deep structure, etc. And I assumed that binding theory properties held only at surface structure, as universal conditions, innate, formulated in UG.

You seem to find that problematic. I do not understand why. It is completely straightforward. Next, you say you find confusing my free usage of such notions as “computational system” while I reject the computer model. Again, no problem at all. The computer model, as it has been developed in the cognitive sciences, doubtless makes use of computational systems—in the wrong way, I have argued. Maybe I’m right, maybe wrong, but there is not the slightest problem. I was using computational systems long before philosophers and cognitive scientists developed the computer models that I have criticized. By a computational system I mean exactly what is meant in the mathematical theory of computation (algorithms, Turing machines, recursive function theory, etc.). I do not see what the difficulty is.

As for your further questions in this connection, they are so remote from anything I have said, or anything even suggested by what I have said, that I am unable to respond. There must be some real failure of communication here, and I’m at a loss to identify it.

Place (December 29, 1993)

It is clear from what you say that I had completely misconstrued your distinction between deep and surface structure. I had supposed that deep structure is a structure in the brains of language users whose existence explains those aspects of the rules governing sentence construction which are linguistic universals in the sense that they are found in every natural language. On this usage, surface structure would simply be the structure whose existence constrains those aspects of the rules which vary from one natural language or group of natural languages to another. Obviously, if this were how the terms were being used, it would only make sense to deny, as you now do, the existence of surface structure, if you were prepared to deny, as plainly you are not, that there is *any* structure in the brains of language users

which gives them the ability to construct well-formed sentences. Given that you don’t want to deny a role for brain structure in linguistic competence, you could only deny that there was such a thing as deep structure if you thought, as plainly you do not, that there are *no* linguistic universals.

From what you say, it now appears that the concepts of deep and surface structure are to be understood in terms of the particular role they play in a complex computational theory of language interpretation and production which I confess I do not understand sufficiently for me to be able to grasp what considerations would lead one either to assert or deny the existence of such structures.

However, what your letter does make clear is that I was right in thinking that for you the criterion for deciding whether or not an aspect of language is innate is whether or not it is a linguistic universal. This is apparent from your discussion of the case and question-formation systems where you assign the systems themselves to C (the computational system) which you take to be innate, while the way those systems feed what you call “the articulatory-perceptual system” varies from language to language and thus shows that “the choice of parameter values” within the innate framework constituted by C is not itself innate.

Chomsky (January 18, 1994)

The terms have always been used as I described, within the discipline itself. Outside, they have taken on all sorts of odd connotations, with “deep structure” commonly used in the sense of “innate structure” or something like that—one reason I abandoned the terms years ago, realizing that there was no hope of putting an end to the free associations that were rampant. Keeping to the technical sense (the only one ever used within the discipline), at the moment, it seems that there is good reason to reject both deep and surface structure, meanwhile con-

siderably deepening and extending the hypotheses about innate determination of languages. I gather from what you say that you aren't interested in the reason for my (current) belief that deep and surface structure can be dispensed with; they'd take more than a few phrases to explain anyway. There's a discussion in some recent technical papers, if you would like to follow it.

You suggest that the criterion for deciding whether an aspect of language is innate is whether it is a universal. That's not exactly false, but that is a highly misleading and confusing way to put the matter. First, recall that a "universal" need not be manifested in every language, as the term is used in its technical sense. Second, that's simply the wrong way to approach the question. I apologize if I'm repeating, but from a naturalistic point of view, the story is (very briefly) like this. There is strong evidence that one component of the brain is dedicated to language and its use: call it the language faculty. For each individual, it has a genetically determined initial state; these are similar enough across the species, apart from extreme pathology, so that we can sensibly speak of the common initial state of the language faculty. In the course of maturation and interaction with the environment, that state undergoes some modification, finally stabilizing (pretty much), apparently around puberty. The states attained can be called "(I-) languages," [internalized languages] in the sense I've explained elsewhere [e.g., Chomsky, 1986]. The theory of the initial state is commonly called "universal grammar"; the theories of the states attained can be called "grammars" of the (I-) languages. We can, if we like, call some of the properties of the initial state "linguistic universals."

A currently plausible theory is that the initial state specifies an invariant computational procedure *C* that maps an array of lexical choices to LF and PF representations, which constitute "instructions" for interface systems, and enter into performance in this way.

C is not entirely invariant; there is some flexibility towards the peripheral end of the component that yields phonetic (PF) representations. As for the lexicon, it is somewhat parametrized: there are options among the formal elements, and of course in the trivial matter of Saussurean arbitrariness. Possibly elsewhere, though it is not clear.

You say that the choice of parameter values within the innate framework is not innate. That's a bit ambiguous. The range of possible choices is innate; the selection among them results from experience. That's why English is not Swahili. There is nothing conceptually obscure here, though the empirical and theoretical problems (which we are not discussing) are extremely difficult and challenging.

Place (May 16, 1994)

As you rightly suspect, I am much more interested in the aspect of language you think is innate and in your reasons for drawing that conclusion than in trying to understand a distinction that you no longer consider to be important in this or any other regard.

In responding to my suggestion that the criterion you are using to distinguish what is innate in language from what is learned, is whether or not the feature in question is a linguistic universal, you seem to be saying that what is wrong with this way of construing the matter is that it assumes a bottom-up rather than, what you prefer, a top-down approach to language. In other words it suggests that the right way to proceed is to survey as wide a variety of different natural languages as possible, identify the linguistic universals and then by induction from those linguistic universals, construct a picture of the innate core of language. As something of a dyed-in-the-wool empiricist, that is the kind of strategy I would be most happy with. However, before conceding that what is linguistically universal must be innate, I would want to see whether this core of

linguistic universals could not be accounted for in terms of what is common to all human natural and cultural environments on the one hand, and the principles of all learning on the other.

In contrast to this bottom-up approach which I would favor, you present me with a top-down a priori theoretical structure which simply assumes that there is an innate language faculty and a universal grammar and that the developmental process whereby full adult linguistic competence is ultimately acquired, is a process of maturation rather than, as I would think, a process of learning. You say that “there is strong evidence that one component of the brain is dedicated to language and its use”; but you don’t say *what* evidence. It seems to me, as I have indicated in a previous letter, that there is almost as little *neurological* evidence for a language faculty in the brain, as there is for a localized memory store. Such evidence as there is, is confined to that for Broca’s and Wernicke’s areas in the dominant hemisphere of the cerebral cortex which, as you point out in your letter, is not nearly as clear cut as is sometimes supposed, or, if I may say so, as your hypothesis would lead one to expect.

I don’t know what other evidence you have in mind; but if, as I suspect, it is evidential confirmation of deductions from the theory which you base on this initial assumption, I would be extremely skeptical for the reasons which Popper has expounded.

Suggested Readings

For Chomsky’s exposition of the deep versus surface structure distinction, see *Aspects of the Theory of Syntax* (Chomsky, 1965, p. 136, p. 198, p. 199, pp. 128–147). For a concise, easy-to-understand account of this distinction, see Pinker (1994, pp. 120–124, p. 475, p. 482). Finally, see Matthews (1993, pp. 185–187) for a discussion of some apparent inconsistencies in Chomsky’s use of the term “deep” in

“deep structure.” According to Matthews, given one set of quotations, one can characterize Chomsky’s more recent theorizing (e.g., the replacement of deep structure by other theoretical entities) as “an elaboration or unfolding of ideas” present in embryonic form in the 1950s. However, given another set of quotations, “one can argue that on the contrary he has changed his mind repeatedly and has sought to conceal it by what amounts at best to rhetorical opportunism” (p. 188).

COMPUTERS AND THE BRAIN

Chomsky (February 24, 1993)

[In “Eliminative Connectionism” (Place, 1992)] you’re quite right in saying that I was not influenced by computer models. In fact, I was skeptical then (and so remain) about invoking them in cognitive psychology (e.g., in the various functionalist approaches, etc.). But the congruence you mention is not remarkable. I was much influenced by the clarification of the concept of recursive (generative) procedure in the formal sciences, which made it possible to formulate traditional ideas and insights about language in a precise way, and to open up vast new areas to empirical inquiry and theoretical explanation. These same ideas show up in automata theory and computer science (and, incidentally, in connectionist models).

[In “Eliminative Connectionism” (Place, 1992)] you say that there is now a plausible alternative to the S-D [serial-digital] computer as a model of how the mind works. Since I never thought of that as a particularly useful model, I can’t comment. However, the idea that work of the PDP [parallel distributed processor] group has suggested a new, or any model, of how the brain works is surely untrue, at least by the standards of the sciences. There is, by now, evidence that connectionist systems of various kinds may play some role in peripheral areas of sensory processing, perhaps with regard to such matters as hyperacuity. Attempts

to apply them to even the most trivial questions of neural function have been a gross failure. The most obvious question, of course, was whether one could model the behavior of nematodes in these terms, given that the “wiring diagram” for this 800-neuron organism is completely known, as is its developmental program. That was indeed tried, but quickly abandoned, because connectionist models abstracted so far from the physical reality that they appeared to be completely valueless. If you are interested, you might contact Charles Rockland, of the Laboratory for Information and Decision Systems at MIT, for a detailed study of his on this topic (only a few paragraphs on connectionism, because of the immediate failures). A review of a variety of efforts concerning language by Lackner and Bever showed that insofar as they succeeded at all, it was because they built in properties of the systems being modeled, though these decisions were completely arbitrary from the point of view of connectionist architecture.

It’s conceivable that something of scientific merit with regard to language and thought might develop from these ideas, despite the record of total failure so far. Similarly, it is conceivable that unstructured systems of unknown properties might some day replace the complex constructions of embryologists in terms of instructions for generation of proteins under particular concentrations of chemicals, etc. But if someone were to make that proposal to embryologists, they would not even bother to laugh. At least by scientific standards, the reaction should be the same with regard to language. It is rather intriguing that it is not—a reflection, in my opinion, of the profound irrationality with which questions of language and thought have historically been treated, and the great difficulty of convincing people to approach them by the methods and standards of the sciences.

I quite agree with . . . [your comments in “Eliminative Connectionism”

(Place, 1992)] about “the predominant fashion in contemporary artificial intelligence”—to which we should, in my opinion, accord the same respect as earlier fashions in AI [artificial intelligence] (the term “fashion” is quite well-chosen for a field that is largely PR hype). The predominant fashions in AI have seemed to me to be without merit, at least as an attempted contribution to science, for 40 years. Its leading practitioners continue to proclaim, as they have been doing for over 40 years, that if you throw a big enough and fast enough machine at some problem you don’t understand, then maybe something will happen. Maybe, but no scientists are holding their breath. Note that I’m making no criticism at all of such matters as expert systems, perhaps quite a useful device for particular purposes. And some real scientists, notably Dave Marr, have made excellent use of the technical achievements in AI to deal with real scientific problems—while, at the same time, denouncing the enterprise in far harsher terms than mine.

Place (June 24, 1993)

I suppose I have to take your word for it that you were never enamored of the S-D computer model as a theory of how the brain works; but you cannot surely deny that others have been understandably impressed by the analogy between your concept of syntax-generating rules in the brain and the rules that make up the programs required to generate computational processes in such machines. As to the effectiveness of connectionist networks in providing a plausible model for the functioning of the brain, I am obviously much more impressed than you are by the ability of such networks to replicate, in particular, the phenomenon of pattern discrimination learning as it is observed in the behavior of prelinguistic organisms, such as the rat and pigeon. With regard to language, although I don’t think that the connectionist attempts to model features of language

learning in the child have been quite such abysmal failures as you suggest, I am inclined to think that such attempts are premature. Much more attention needs to be paid to ensuring that both the microstructure and functional properties of connectionist networks correspond to those of the brain, that the learning capacities of prelinguistic organisms are accurately modeled, and that the actual learning conditions experienced by the language-learning child are closely followed, before any strong claims are made about the ability of such devices to replicate the acquisition of linguistic competence in the form in which it actually takes place in the real world. Where you and I differ is in our respective prognostications of how things are likely to turn out when such work is eventually done. In that regard, I would remark only that someone who thinks that something can and eventually will be done is always in a stronger position, in that his view is less exposed to falsification by the way things turn out, than someone who denies that such an enterprise can ever succeed. Of course, you don't in fact do that in your letter. You simply suggest that to propose that such enterprise might succeed is so ludicrous as not to be worth taking seriously. To that all I can say is *chacun à son goût*, while recognizing of course, that your *goût* carries a lot more clout in the academic world generally than does mine.

Chomsky (July 22, 1993)

Doubtless many people have been enamored of such models. I never have, as the literature makes clear enough. Computer models may be informative here and there; they are often extremely misleading. Many of the topics that engage philosophers and philosophically minded cognitive scientists (e.g., Chinese room, "the computer model of the mind," etc.) seem to me radically misconceived from the outset, for reasons I've written about.

As for connectionism, there is, at the

moment, about as much reason to take connectionist models seriously for language as for embryological development. True, it's not excluded that they might somehow do something, in either case. So far, they are utter failures, except at the very periphery; there is some plausibility to the resort to such ideas for peripheral processing, as Tom Bever and others have suggested. The issues, incidentally, seem to be becoming moot. Connectionist language modelers seem to have largely gone off to other things, perhaps convinced by repeated failures to achieve any results. As for connectionist models of the brain, as I may have mentioned, they seem inadequate even for the case of nematodes, with 800 neurons and a wiring diagram and developmental pattern completely understood. I think . . . on matters such as the brain and connectionism, . . . [Gerry Edelman] seems to me much on the mark. I doubt that many brain scientists would disagree.

It may be, as you say, that we differ on prognostications as to how things will come out, but that seems to me a bit misleading. Similarly, if someone suggested that connectionist models could replace embryology, I would disagree with the prognostication, but to say just that would, again, be misleading. More crucially, we differ on the point, or pointlessness, of making claims about—or even toying with—ideas for which there is no evidence at all, and which appear to have little intrinsic interest. If someone said that the brain will be explained some day by electromagnetic theory, they could be right, but I wouldn't spend much time thinking about it until some plausible reasons are adduced. Same with connectionism.

As for whose *goût* carries more weight in the academic field, mine carries very little. Ask your friend David Armstrong, who you mention. Count noses among people who speculate about these matters (philosophers, cognitive scientists, etc.). The virtues of connectionism win hands down—

more's the pity. Few are even aware of my views on the topic, and what I write, they don't read.

Place (August 21, 1993)

With regard to traditional computer models, we agree. I would also agree that current connectionist models of language do not get us very far. But the reason for that, as I see it, is they have tried to run before they can walk. What they should have done, and what some of them are beginning to do, is to see how far a connectionist network can reproduce the phenomena of simple pattern discrimination and expectation learning as it is observed in animals. Only when that foundation has been properly secured, can we expect genuine progress towards modeling the complex process of language learning in the child as revealed by studies such as those of Ernst Moerk (1983).

With regard to connectionist networks as models for the brain, having spent three months in his Neurosciences Institute in New York in 1991 and having read . . . much of *Neural Darwinism* [Edelman, 1987] . . ., I fully agree with your assessment of Gerry Edelman . . . with respect to the soundness of his criticism of current connectionist modeling. But what Gerry is criticizing is the connectionist's lack of attention to what is known from neuroscience about the way the brain is actually organized and actually operates. In particular he criticizes (Reeke & Edelman, 1988) the principle of learning by the back-propagation of error-correction on the grounds that all known causal effects within the brain feed forward from input to output and never in the reverse direction, except in so far as excitation from the output end of a network is fed back into the input end by means of what he (Gerry) calls "re-entrant circuits." What impresses me about Gerry's work is not so much the way it differs from connectionist modeling, but the things the two approaches have in common. They agree in holding (a) that the brain con-

sists of a complex network of synaptically interconnected neurons, (b) that the crucial factor in determining the properties of the network as a whole is the number and "weight" of the individual synaptic connections between the neurons of which the network consists, (c) that learning is effected by selectively increasing and decreasing the weights of individual synaptic connections, (d) that nothing in the brain is "hard-wired"—though the possibilities of further modification contract with age, no connections are completely impervious to change, none are unaffected by previous activation, (e) that changes in synaptic weights come about in accordance with what connectionists call "learning rules," (f) that connectionists, following Jordan (1986) are increasingly emphasizing the role of what they call "recurrent circuits" which are indistinguishable from Edelman's (1987) "re-entrant circuits." It also seems to me that what Gerry talks about as "categorization," which is evidently a form of prelinguistic concept formation, is essentially the same process as the pattern discrimination learning observed both in animals (e.g., Lashley, 1930, 1938; Herrnstein, Loveland, & Cable, 1976; Pearce, 1988, 1989) and in connectionist networks (Churchland, 1988, pp. 157–162).

Chomsky (November 8, 1993)

As for connectionism, it is conceivable that there is something to these particular abstract mathematical models, but for the moment, their interest (at least, for language) appears to be approximately zero. As perhaps I've mentioned, they were also quickly abandoned, as abstracting much too far and in the wrong ways from physical reality, in the study of nematodes. In the case of language, there seems to be nothing more to discuss, because, as far as I am aware, the entire enterprise has been abandoned after its failures (accompanied by much hype, as is the usual story in areas around AI).

As for Edelman, doubtless he is a fine biologist. I have no criticisms of his work. It simply has no relevance to these topics, as far as he has shown.

Place (December 29, 1993)

While this exposition clarifies your position very considerably, it does not bring us any closer to agreement. For although you claim to reject the "computer model" your conception of what you call "the language faculty" is essentially computational. From my perspective, the fact that your theory owes more to the "mathematical theory of computation" than to any actually constructed hardware is beside the point. Your theory assumes that the process whereby a language-user interprets or constructs a sentence is a computational process, a process in which a function is computed in accordance with an algorithm and where computation is essentially a matter of reading and manipulating symbols in accordance with a rule (the algorithm). It assumes moreover, that "the language faculty" is an anatomically distinct box of which it makes sense to say that it "interfaces" with other equally discrete boxes, the "articulatory-perceptual" and "conceptual-intentional." It also postulates the subdivision of the language-faculty box into two subboxes, the "computational system" which, I take it, handles syntax and "a lexicon" which presumably, since it apparently predates the choice as to whether a concept is labeled "tree" or "Baum" contains a collection of (?) innate concepts or meanings to which such labels are subsequently attached. To my mind, this story is pure fiction. It is useful, no doubt, in drawing attention to the different aspects of linguistic competence and the way they are related to one another, but as an account of the actual states of the brain underlying the ability to construct and construe novel sentences, it is no more credible than earlier attempts to map nominalized psychological predicates onto the brain, such as that of phrenology.

I have two principal objections to this scheme, one is to the idea that there is a language faculty which constitutes a distinct unit in the brain with no other functions, the other is to the idea that the prelinguistic brain is in the business of symbol-manipulation. We know of course, that there are two areas in the human cerebral cortex which are specialized for linguistic functions, Wernicke's area and Broca's area. But the functions of these two areas are quite specific. Wernicke's area is concerned with the ability to supply lexical words as labels for concepts where the concepts are generated and selected elsewhere. Broca's area is concerned with the detailed syntactic organization of sentences whose underlying thought is again generated elsewhere. The inescapable conclusion is that the function of these areas is to give a specifically linguistic form to the products of a behavior-organizing system which long antedates the evolution of language and occupies the remaining 99%, or whatever the correct figure is, of brain tissue.

My other objection is to the claim that the brain is a computer in the sense of Alan Turing's mathematical definition according to which a computer is a device which, given the appropriate algorithm and enough time, can in principle compute any function. My difficulty with this claim is brought out by another paper which I heard at the conference in Slovenia which I mentioned in a previous letter. It was given by a philosopher from München University by the name of Gerhard Helm (1993). Taking Turing's definition as his starting point, he generated the following proof:

1st premise: Given the appropriate algorithms, and enough pencil and paper, I (i.e., Gerhard Helm) can in principle compute any function.

2nd premise: By Turing's definition a computer is something/anything that given the appropriate algorithms and enough time can compute any function.

Ergo, I (Gerhard Helm) am a computer.

3rd premise: I can think.

Ergo, at least one computer can think.

Now the thing to notice about this argument is that the computer is Gerhard Helm, not Gerhard Helm's brain. Gerhard Helm only qualifies as a computer when equipped with the pencil and paper required to do his calculations with and on. What this tells us is that the human brain only becomes involved in the kind of symbol-manipulation in which computation in the sense of Turing's definition consists, when it not only already has the ability to communicate with others by means of natural language, but has also acquired the ability to write and thus commit its symbol manipulation to paper or other writing material, whether it be papyrus, parchment, or tablets of wax or clay. Before that stage is reached, the brain is what the late Donald Broadbent (1991) has called "a subsymbolic processor," in other words, a connectionist network.

Chomsky (January 18, 1994)

In the philosophy-cog.sci. literature there is something called the "computer model" of the brain, outlined in various functionalist approaches (see, e.g., Ned Block's account [Block, 1990] in the third volume of the Osherson et al. *Invitation to Cognitive Science* for a lucid and typical account). I do not accept any of this.

However, there is good evidence that the states and properties of the language faculty are accurately described in computational terms. That is, one property of the initial state is that it incorporates an (almost) invariant computational system that generates (PF, LF) pairs, etc.—or so evidence now seems to show. This is all simple straightforward normal science, exploring states and properties of the brain and seeking empirical evidence for theories about them. There are no philosophical problems, apart from the kind that arise in any study of states and properties of objects in the world.

You say that the picture is "pure fiction"; fine, as long as you recognize that your reaction is pure prejudice.

Thus you "confess (you) do not understand sufficiently" the theories that you are dismissing as fiction, which is, in my view, an odd way to proceed. If you think there is something wrong with trying to show that a complex system has states and properties, then it would be interesting to hear the reasons. As you state explicitly, you are not in a position to say anything about the particular hypotheses about these states and properties—except that they are fiction, like phrenology. At this point we've degenerated to complete unreason, as far as I can see. It's as if some biologist were to tell me that he thinks the states and properties of the cell are so-and-so, and I'd respond by saying that "I confess I do not understand his theories sufficiently for me to grasp what would lead one either to assert or deny the existence of such structures"—but I do know that it is "pure fiction," on a par with phrenology. If he bothered to respond he'd say, I guess, that I should make the effort to understand the theories before passing judgment on them.

You say you object to the idea that there is a language faculty that is a part of the brain with no other functions. I quite agree with your objection. Thus, there is a component of the body (the brain) that is dedicated to thought, language, etc. But it surely has other functions. Thus, it seems to be a crucial part of the circulatory system. You further say that "we know of course that there are two areas specialized for linguistic functions." We "of course" know nothing of the sort. Even identification of Wernicke's area is quite a dubious prospect, as Joseph Bogen pointed out some years ago. What is known about the brain is so rudimentary and primitive that it contributes almost nothing to the questions we are discussing. It's kind of as if a physicist a century ago had said that "we know of course that the physical world involves mechanical and electromagnetic forces," and they therefore account for the chemical bond, the properties of the elements, etc. That was just untrue;

physics turned out to be radically false. The brain sciences are nowhere within shouting distance of 19th (or 17th) century physics. We simply don't know what aspects of the brain are implicated in language, or even whether brain scientists are looking in the right place.

The reference to Turing machines is entirely beside the point. No one suggests that the computational procedures are arbitrary general recursive functions. You might as well object to physics, on the ground that unstructured theories with no specific hypotheses are unlikely to explain anything; true, but is it worth saying?

Place (May 16, 1994)

It is clear from what you say on this score that the computer model you repudiate is a subspecies of what I would regard as computer models in general, of which your own theory is another subspecies. I want to repudiate all such computer models, including yours. What I find unacceptable is the notion of rules playing a generative role in relation to language and behavior. Rules for me play a normative, error-correcting role with respect to processes which are generated by irrational associative principles. It is because this is how a connectionist network operates that I find connectionism appealing.

With regard to my claim that your language faculty is a fiction, I must acknowledge that this is derived in part from a long-standing prejudice against believing in the existence of *any* kind of abstract object of which the faculties of the mind are just one instance. However, this is not pure prejudice, as you suggest, since it is based on what I regard as a sound linguistic argument, though again this is not a view that I would expect you to share.

The argument can be illustrated by means of the sentence *John gave Mary the book*. This, as I think you will agree, consists of a predicate in the form of the past tense of the verb *give*

which creates three argument places, the giver, the receiver, and the object given, all of which are occupied by noun phrases denoting what Aristotle calls "substances," in this case *John*, *Mary*, and *the book*. As you will also agree, there are two active-passive transformations of the sentence, namely, (a) *Mary was given the book by John*, and (b) *The book was given to Mary by John*. These transformations do not change the semantic content or the truth value of the sentence in any way. However, by putting each occupant of the three argument places in turn into the all-important subject position, they have the effect of altering the point of view from which the event in question is viewed. Thus *John gave Mary the book* looks at the event as action on John's part and to that extent from his point of view. *Mary was given the book by John* looks at the event as something that happened to Mary and hence from her point of view, while *The book was given to Mary by John* looks at it as something that happened to the book and hence from the point of view of someone interested in its history.

But there is also another transformation in which it is the predicate that goes into the subject position, as in the phrase *John's gift of the book to Mary*. This differs from the other transformations in that it is a noun phrase rather than a complete sentence, one which not only focuses attention on the event rather than its participants, but has the function of permitting the construction of a sentence in which the event denoted by the predicate in the original sentence, occupies an argument place relative to a second order predicate, as in the sentence *John's gift of the book to Mary was extremely generous* or *John's gift of the book to Mary made Joe's gift of a pencil look mean*.

In order to achieve this transformation, the original predicate, the verb *gave*, has to be nominalized, that is to say converted into a noun, in this case the noun *gift*, and it is these nominalizations of predicates and other non-

substance denoting parts of speech which, according to me, are the source of fictitious abstract objects. In the case of the noun *gift* the temptation to suppose that this denotes an abstract object over and above those concrete objects, the giver, the receiver, and object given, is minimal. This is partly because what is referred to is a particular event rather than a type of event, and partly because the specification of the concrete objects occupying the original argument place makes its derivation from the original sentence very clear. But when we begin to talk in generalities about such things as memory, perception, or language, we lose the connection with sentences about people remembering or recognizing things, seeing, hearing, smelling, tasting, and feeling things, saying something, speaking, talking, writing, listening, understanding, and reading what is said or written.

If that is how abstract objects such as the faculties of the mind get generated in the first place, it is hardly surprising that they fail conspicuously to line up with anything that neurology tells about the way psychological abilities break down when different parts of the brain are damaged. It is not just phrenology which illustrates this, the total failure of neurology to come up with anything resembling the long-term memory store of cognitive psychology is just as striking; though in this case the notion persists despite the evidence against it, because of the pervasive influence of the computational model with its information-storing disks and tapes.

Clearly I take a more sanguine view than you do of the current state of neuroscience. I also think that comparisons with physics are not particularly helpful when considering the state and likely future development of the biological sciences in general and neuroscience in particular. But here we are back in the domain of hunch and *goût* rather than solid ground of argument and evidence.

Suggested Readings

See Chomsky (1993, pp. 29–31, pp. 42–44, pp. 90–93) for some brief but elucidating comments on his position regarding computer modeling of the mind. For Place's view, see Place (1992). Also, see Block (1990).

POP-DARWINIAN FAIRY TALES

Chomsky (February 24, 1993)

[In "Eliminative Connectionism" (Place, 1992)] you say that a "serious objection" is that I provide no account of how such a complex and sophisticated piece of genetic endowment as linguistic theory postulates "might have evolved." That's a strange comment. The complexity and sophistication are slight, as compared with what is postulated in embryology, for example, and no one has any idea how these systems might have evolved, or indeed if they evolved at all. The problem only falls more out of control when we consider postnatal aspects of maturation and development that are invariably assumed to be genetically determined, though no one has any idea how these complex and sophisticated developments take place: onset of puberty, for example.

At least for serious evolutionary biologists, the comment you make would appear virtually meaningless. I'd suggest, for example, that you have a look at Dick Lewontin's chapter on "The Evolution of Cognition" in Osherson and Smith's *Invitation to Cognitive Science* [Lewontin, 1990]. Commenting on various misconceptions of evolutionary theory, he writes,

First, there may have been no direct natural selection for cognitive ability [or language] at all. Human cognition may have developed as the purely epiphenomenal consequence of the major increase in brain size, which, in turn, may have been selected for quite other reasons.

(Perhaps thermoregulation, he suggests, possibly tongue-in-cheek.) If so, it would be like much else in the biological world.

As for “the same basic learning capacities as are available to other mammals,” what we know about other mammals is that they have quite a variety of cognitive structures, highly adapted to particular problem spaces, and the question of what in these developments might constitute “learning” is, for the moment, very obscure, to my knowledge. Indeed, it is not very clear that the concept itself belongs in scientific psychology, except possibly in some very peripheral role.

Place (June 24, 1993)

You cite systems “postulated in embryology” as a far more complex piece of genetic endowment than that postulated by your version of linguistic theory and one the manner of whose evolution is equally obscure. What you fail to notice is that the “systems postulated by embryology,” assuming they have been correctly described, have had a vastly longer period in which to evolve than those postulated by your version of linguistic theory. I can’t say that I find Lewontin’s version of magical emergentism which you quote particularly convincing. There are undoubtedly cases in the evolutionary story where features which have been selected by virtue of possessing one kind of functional utility have been further developed by virtue of subserving some other and quite different function. But I fail to see any convincing reasons to suppose that the evolution of the innate linguistic endowment postulated by your theory is a case in point. In the necessary absence of direct observational evidence, this is a case where the issue is going to be decided by an application of Ockham’s razor. If, as I suggested in my paper, the development of language can be explained in terms of the principles of animal learning theory, when combined with the special circumstances of [an] ape with highly developed manipulative abilities, a finely tuned voice box and palate and the increased brain size required to exploit those manipu-

lative and vocal skills to the full, an innate cognitive ability specialized for the generation of syntactically well-formed sentences would be convincingly “occamed,” to use my friend, David Armstrong’s, apt phrase.

As for your doubts about the place of the concept of learning in scientific psychology, it is undoubtedly the case that the so-called “cognitive revolution” in psychology for which your work in linguistics was a major stimulus has sidelined the issue of learning which had been central to psychology during the period when behaviorism was still a fashionable view.

What for me is enormously and heart-warmingly exciting about the connectionist counterrevolution which is now well under way, is that it restores the concept of learning to what those of us who were never seduced by the S-D computational model see as its rightful place as the central concept in a science of psychology, a science which focuses, in the way recent cognitive psychology has failed to do, on what is common to the behavior of all complex free-moving living organisms, in contrast to the species-specific differences which preoccupy the ethologist.

Chomsky (July 22, 1993)

First, Lewontin is not advocating “magical emergentism.” He has arguments, and they are not so frivolously dismissed. He is making the simple and accurate point that where we lack understanding, there is no point constructing fairy tales, though it is easy to do so, one way or another. And it is a truism that there are many factors, mostly quite poorly understood, that enter into determining the form of complex systems, including organisms; biophysical laws, for example.

You say that I offer no “convincing reasons” that the evolution of the innate language faculty is a case of development by virtue of subserving other functions. Absolutely correct; I’ve nowhere suggested anything different.

Exactly Lewontin's point as well. There are no convincing reasons for anything in this domain, which is why it is pointless to propose vacuous natural selection fairy tales, as Lewontin observes. Note that I'm not offering a different fairy tale, as you seem to have misunderstood me to be; rather, eschewing fairy tales.

As for the development of language being explainable in terms of animal learning theory, etc., that can be true only in a sense of "explain" that we should not take seriously. It is not quite accurate that I "failed to notice" the differences of evolutionary time scale. Rather, I failed to mention them, because they are entirely irrelevant to this discussion, in the present state of understanding.

Place (August 21, 1993)

I cannot agree with your dismissal of attempts to reconstruct the evolution of language as mere "fairy tales." It is, of course, true that since spoken language leaves no direct traces in the fossil record, we do not have and cannot expect to have any direct evidence bearing on this issue. To that extent any attempt at reconstructing the process whereby human natural language evolved is inevitably a matter of speculation, rather than observation. Nevertheless Darwin's principle of variation and natural selection has proved itself time and again as the principle governing not only the evolution of biological species, but, as Gerry Edelman among others has pointed out, other developmental processes in the domain of the biological and social sciences. It has even been suggested (Gehrz, Black, & Solomon, 1984) that it has application on a cosmic scale to the evolution of galaxies, stars, and planetary systems. Given that principle, it is not difficult to sort what you are calling "fairy tales" about the evolution of language into those which are plausible, in that they conform to Darwin's principle, and those that are implausible because they do not.

Chomsky (November 8, 1993)

If I said that stories about language evolution are "fairy tales," that was a bit too strong, though not much. There are some interesting speculations (e.g., Jerison and others) about possible connections between language and imaging areas, and a few other things have been suggested. But in my opinion, Lewontin is basically correct about the prospects (given current understanding), and I'm aware of nothing that even begins to face the most elementary questions. Much of the work . . . is just confused, in my opinion, for reasons that I and others have discussed. As for "evolution of biological species," as far as I can see little is known. Not that Darwin is wrong; rather, as he insisted, natural selection is only one factor, and we have little reason now to believe that it is a major factor in determining why insects are what they are, etc. As for stories about evolution of language that conform to Darwin's principle, they are myriad, and one can make up more at will. Furthermore, none have any serious status, apart from the most marginal considerations. It's for this reason that it is quite fair to describe them as fairy tales, like a lot of pop-Darwinism.

Place (December 29, 1993)

It is true that natural selection alone cannot account for the evolution of anything. Before natural selection can get to work there has to be variation and variation has to be explained, and explained independently of the process whereby some of *those* variants are selected in preference to others. But given a story that accounts plausibly for both variation *and* subsequent selection and does so more economically than any rival theory, we can, I suggest, be reasonably confident, at least until a new and better theory comes along, that that is what actually occurred. If what we are comparing here is a nativist and an empiricist theory of the evolution of language, I have personally

no doubt that on all these counts, the empiricist story wins hands down.

Chomsky (January 18, 1994)

Note that I don't use . . . [the phrase "Darwinian fairy tales"]; rather "pop-Darwinian fairy tales," based on utter misunderstanding of evolutionary theory. The basic problem is to determine the space of physical possibilities within which mutation and selection take place. That's a matter that is little understood, and not understanding it, we can say very little about selection in the case of complex systems. These are not points of much dispute among serious biologists. We are not, as you suggest, comparing "nativist" and "empiricist" theories of the evolution of language, so neither wins "hands down." There is a little very preliminary work on the topic that can be called "science," and is quite neutral between what (I guess) you have in mind in speaking of "nativist" and "empiricist" theories. Beyond that, we have fairy tales concocted by people who mistakenly believe that they are being "empiricist." I've never seen anything called a "nativist" theory of the evolution of language. There is some deep confusion in your exposition here, I'm afraid. Beyond a few intriguing speculations about relations between imaging and symbolic communication, and the like, we have two positions: (1) Lewontin's, which I share, that we are in no position to propose any serious account; (2) a variety of fairy tales concocted by people who typically labor under the delusion, anathema to Darwin and dismissed by serious biologists, that all traits are "selected," and who think of themselves as "empiricist": a problem of cultural history, perhaps, but hardly more than that, in my opinion.

Place (May 16, 1994)

I agree completely (a) with your claim that a Darwinian-type explanation must explain the source of the initial variation within a population, as

well as the subsequent process of selection between variants, and (b) with the point that any plausible account of the evolution of a complex phenomenon such as language must invoke *both* the development of innate capacities peculiar to the species *and* cultural learning by the individual, and that, therefore, the contrast between nativist and empiricist theories is a matter of the degree to which a particular theory emphasizes the one rather than the other.

Where we part company is over the issue of whether we have enough respectable evidence to go on to begin constructing such a theory and thus decide how much relative weight to attach to the two factors. You think that any such attempt is premature. I disagree. What impresses me particularly is the recognition by the likes of Gerry Edelman of something which those of us brought up in the tradition that stems from E. L. Thorndike and B. F. Skinner have long been familiar, namely that the principle of variation and natural selection applies just as much to the processes of *ontogenetic* development as it does to the processes of *phylogenetic* development. This, together with emphasis on learning within connectionism and what, in evolutionary terms, is the very short time span over which human language has evolved, strongly suggests that a theory that relies maximally on individual learning and minimally on inherited characteristics is going to be the right one.

Suggested Readings

For some comments on the role of evolution in language, see Chomsky (1980, pp. 99–100, pp. 229–231, pp. 239–241, p. 249; 1988a, pp. 158, pp. 167–170, p. 189; 2000, p. 163). For an easy-to-follow (albeit brief) discussion of Chomsky's views on evolution and language, along with a counterargument, see Pinker (1994, p. 333, pp. 362–364); for a more elaborate treatment, see Pinker and Bloom (1990).

See Place (1992) for some brief comments on his view. See also Lewontin (1990).

BEHAVIORISM

Chomsky (February 24, 1993)

[In "Eliminative Connectionism" (Place, 1992)] you quote a comment of mine [about behaviorism] from a series of lectures to a group of undergraduates at a teacher's training college in Kyoto which reviewed the history and modern development of generative grammar, and you ask why I still find it necessary to devote a few phrases in three lectures to a point of view which, in my opinion, made no sense in the first place and was discredited and largely abandoned (rightly, in my opinion) decades ago. There is no "lurking suspicion that the corpse may yet rise up." Rather, it seems that in this context, a position that (unfortunately) dominated the study of psychology for many very barren decades merits a few phrases. There is much more discussion in the same lectures of structural linguistics, which does not have a lot more to recommend it than behaviorism, in my opinion. Again, it would seem quite improper, in that context, to omit it entirely. I don't recall whether I also referred to Piaget. I should have, if I didn't, again without any lurking suspicions. As for the "new way of looking at how the mind/brain operates . . ." I'm afraid that is just hand waving, for now.

Place (June 24, 1993)

You complain that, in the rhetorical flourish which concludes my . . . ["Eliminative Connectionism" paper (Place, 1992)], I took your remarks about behaviorism in your Kyoto lecture out of context, with the result that I made them seem to attach more importance to refuting behaviorism than to refuting other bankrupt approaches to language, such as structural linguistics and Piaget, which were also mentioned. Part of the reason for this is that

I came across this passage in a complimentary copy of Lycan's *Mind and Cognition* [Lycan, 1990] which I received as a consequence of the inclusion in that book of my own paper "Is Consciousness a Brain Process?" [Place, 1956]. The passage quoted by Lycan does not include the references to structural linguistics and, assuming he too was mentioned, Piaget. But this editing on Lycan's part, even if it misrepresents your personal view of the relative importance of previous linguistic theories, does surely capture the central place that behaviorism has occupied in the demonology of linguistics and cognitive psychology ever since it was put there by your 1959 review [Chomsky, 1959] of Skinner's book *Verbal Behavior* [Skinner, 1957]. I have lost count of the number of times over the past thirty years that I have had to listen to papers by linguists and cognitive psychologists which begin with a ritual denunciation of the evils and bankruptcy of behaviorism from which, so we are told, we are now thankfully delivered. And it still goes on. I heard another example only last week in a paper by an American computationalist philosopher at a conference in Slovenia, on "Connectionism and the Philosophy of Mind." You may not fear the resurrection of the corpse; but there's too much protestation about [it] to convince me that others in your camp aren't getting rattled.

Chomsky (July 22, 1993)

If you think . . . [behaviorism occupies a place in the demonology of linguistics and cognitive psychology since my review of Skinner], then by all means write about it, citing sources, which should be easy if the demonology is as passionate as you allege. But please leave me out of it.

I'm responsible for what I write, say, teach, etc., not for the demons of others. For me, behaviorism is a topic of almost no interest. I wrote about it in 1957 (the Skinner review) because of what seemed to me a pernicious influ-

ence that ridiculous ideas were acquiring, and still have, over a very broad range, including work influenced by Quine. I then dropped the matter, with the exception of occasional comments here and there, and a review [Chomsky, 1973, pp. 104–150] of Skinner's *Beyond Freedom and Dignity* [Skinner, 1971] as part of a more general discussion of "psychology and ideology." Behaviorism is not part of my demonology; rather, in my opinion, a curious deviation from rationality and science that merits explanation as a chapter of intellectual history, but one that I haven't spent a lot of time on.

As for linguistics and cognitive psychology, I don't see the "demonology" to which you refer. Rather, behaviorism was simply abandoned, quite quickly, as the absurdity it was. In linguistics, there's been nothing much said about it, to my knowledge. And precious little in cognitive psychology either, I believe, though I know the field(s) less thoroughly.

You say that you've heard countless papers by linguists and cognitive psychologists that open by denouncing behaviorism. I can't recall having heard any, apart from one context. When I'm asked to give a general talk discussing current work and how it evolved, I often start with some comments on the intellectual climate of the 1950s, when behaviorism was indeed taken very seriously. Maybe these comments would sound to you like "denunciations," though to me they sound like descriptions of the intellectual environment of the day. I suppose others do the same, in such circumstances, as they should. But that's a far cry from what you describe. I'm not incidentally trying to deny your experiences. I just haven't had them. Again, I'm not suggesting that you not write about what I haven't experienced and am unaware of: rather, that you give sources and not make blanket charges. That seems fair enough.

Note that I'm not speaking here of what is called "cognitive science" by philosophers, AI people, and some

psychologists. But that's another matter. Your example is "an American computationalist philosopher," who either was a "rattler" or "rattled," I didn't quite follow which. Whatever it was, note that I quite agree that computationalist philosophers are much exercised over these issues, wrongly in my opinion.

As for my "camp," if you are referring to people with whom I work on topics of language, philosophy of mind, cognitive psychology, etc., they surely are not trembling about the resurrection of the corpse, as you suggest. In fact, they pay very little attention to the matter, as far as I am aware. Speaking only for myself, I teach an introductory graduate course on language and mind for linguists, philosophers, psychologists, and many others. There may be a few sentences about behaviorism, hardly more, because no one is interested (including me). I certainly keep an eye on work in connectionism, and when anything appears that seems possibly relevant to the kinds of questions that have interested me, I follow it. That hasn't been an onerous task, so far, but if something comes along that seems worthy of attention, I'll gladly spend more time on the matter, not being "rattled" in the least; rather, curious and eager to learn something, as often in the past.

Place (August 21, 1993)

Your response to my allegation that behaviorism has occupied "a central place . . . in the demonology of linguistics and cognitive psychology ever since it was put there by your 1959 review of Skinner's book *Verbal Behavior*" is to say that whatever may be true of others, behaviorism is not part of any demonology of yours, since to treat it as a pernicious doctrine to be stamped out at all costs would be to accord it more significance and importance than this "curious deviation from rationality and science" deserves.

I don't doubt that that is your view. Nor do I think there is much profit in

arguing about whether or not those who undoubtedly have taken and still take their cue from you in these matters, do or do not show signs of taking the possibility of a resurrection of behaviorism more seriously than you do. What I do think is that before dismissing behaviorism as a palpable absurdity, you should consider what you mean by the term. Behaviorism, as I construe it, is not just a single doctrine or a single project, it is a loosely connected family of doctrines and projects such that commitment to one form of behaviorism need not entail commitment to some or all of the others. Most of these doctrines have a sound rational foundation, though behaviorists have not, on the whole, been very good at demonstrating this to others. Those that are obviously absurd, such as the denial that mental events exist or that all behavior consists of a chain of mechanical reflexes have either never been held by anyone or have long since been abandoned. For me, behaviorism stands for the following principles:

(i) *An objectivist epistemology* (Galileo, Comte, Watson). The observation sentences required in order to anchor empirical knowledge to the reality it depicts are descriptions of relatively permanent states of affairs in public intersubjective space on whose correct description in terms of the relevant natural language or technical code, any competent member of the community of speakers of that language or code who is seriously exposed to that state of affairs will agree. There is no implication that only such states of affairs can be known to exist.

(ii) *Behavior as the subject matter of empirical science* (Wundt, Watson). Every empirical science has and must have as its subject matter the observation, description, and explanation of the behavior of entities of some kind, such that the existence of entities of that kind and the identification of the phenomena being studied as the behavior of those entities is not in doubt.

(iii) *"Mentalism" unacceptable in a scientific explanation of behavior.* For

these purposes "mentalism" is to be defined as the practice of explaining the behavior of a free-moving living organism or animal by quoting what it would say (if it could talk) about the situation confronting it. Such explanations, though defensible up to a point in explaining behavior that is in fact controlled by a linguistic formulation of the contingencies involved, have no place in a scientific account of the behavior of a prelinguistic organism or of the way such an organism acquires linguistic competence.

(iv) *Linguistic behaviorism.* The thesis that linguistic competence is acquired by the child and maintained in the adult human by the same process of "contingency-shaping" as is observed in the acquisition and maintenance of nonverbal skills in a prelinguistic organism.

(v) *Conceptual behaviorism* (Wittgenstein, Ryle). A majority, but only a majority, of our ordinary mental concepts refer either to dispositional states of an organism (usually human) or to instantaneous events whereby such a dispositional state comes into existence. Such dispositional states, though in my view causally dependent on states of the brain microstructure, consist solely in the capacity or tendency to talk and otherwise behave in a range of broadly specifiable ways.

Chomsky (November 8, 1993)

You give several principles that you see as part of behaviorism. Of these, (i) and (ii) have absolutely nothing to do with behaviorism. You're simply describing normal science. (iii) is meaningless, until you give some account of what you mean by "mentalism," which will, I presume, presuppose an account of what you mean by "physicalism" (the physical, body, or whatever). Since no such account has even been remotely hinted at, at least since Newton demolished the mechanical philosophy, we can't talk about (iii). The example you give is irrelevant to (iii), whatever it means (if any-

thing); normal science, again, would disregard “explanations” of the kind you mention, just as it would disregard pop-Darwinian fairy tales about language evolution, even if they fall within the astronomical range of possible accounts that conform to legitimate principles, like natural selection. (iv) is an empirical hypothesis, which seems glaringly false; there is overwhelming evidence against it, none for it that I know of, which is why, to my knowledge, it has largely been abandoned in the study of language acquisition. But true or false, it has nothing to do with behaviorism. (v) also appears to be false, insofar as it is at all clear (which is not very far), for reasons I’ve discussed elsewhere, and won’t repeat.

I agree that the theses you mention are not “obviously absurd.” In fact, most of them are “obviously correct.” They simply have nothing to do with behaviorism, for the most part.

Place (December 29, 1993)

Despite your earlier disclaimers of any sympathy for Kripke’s position, you appear to think that behaviorism is a Kripkean “natural kind” which pre-existed its baptism with that name at the hands of J. B. Watson in 1913. How else am I to interpret your contention that the doctrine of epistemological objectivism, the doctrine that an empirical science should be defined by reference to the objectively observable behavior of the entities it studies, and the doctrine that linguistic competence is acquired by the same processes as are observed in the behavior of prelinguistic organisms (animals and prelinguistic children) have nothing to do with behaviorism?

On any remotely plausible view, behaviorism is constituted by the various doctrines to which those who have called themselves behaviorists have subscribed. If, as most critics of behaviorism do, you characterize behaviorism in terms of doctrines that no one who calls him or herself a behaviorist subscribes to, you simply create a con-

venient whipping boy that fails to correspond to anything in reality.

If you look at the history of behaviorism, you cannot fail to be struck by the fact that J. B. Watson and all his early followers were comparative psychologists who were studying the behavior of animals and young children as a way of throwing light on the mental life and behavior of human adults from within an essentially Darwinian conceptual framework. They studied objectively observable behavior because that is all there is to study in the case of animals and prelinguistic children. They rejected introspection as a source of information, partly because such evidence is not available in the case of animals and prelinguistic children, but also for the same reason that Comte rejected it, because its essentially subjective nature made it unacceptable as the evidential basis for an empirical science. They rejected the notion of psychology as the science of the psyche or mind on the same grounds that Wundt had rejected it earlier, because they thought, as he did, that a science should be defined of what it studies, not in terms of a theoretical construct which may or may not turn out to be useful in explaining what is observed. They rejected Wundt’s alternative definition of psychology as the science of immediate experience on the grounds that immediate experience is not and cannot be the object of study in the case of animals or, indeed, in the case of humans other than oneself. They saw the principle of learning derived from the objective study of animal behavior as providing the key to understanding of linguistic competence, precisely because, unlike the later ethologists, they were interested in animal behavior for the sake of the light they hoped, and for good Darwinian reasons, expected it would throw on the behavior of humans. I don’t see how in the face of this you can possibly claim that objectivism in epistemology, the principle that a science is to be defined in terms of what it studies, and the explanation

of language acquisition in terms of the principles of animal learning are not central to behaviorism as a standpoint in psychology.

With regard to my third behaviorism principle, the issue of mentalism, you appear to assume (a) that I was using mentalism in a sense which implies a contrast between the mental and the physical and (b) that I was using the case where language which implies linguistic competence and linguistic control is used to "explain" behavior where such competence and control is absent as an *example* of mentalism, in some sense in which "the mental" is being contrasted with "the physical." Both these assumptions are false. I have believed in the incoherence of the mental/physical distinction ever since I was persuaded of it by Ryle when I was an undergraduate at Oxford in the late 1940s. It is pleasing to find another philosopher these days who shares that view. They are rare.

When I spoke of the use of language implying linguistic competence and linguistic control when offering a scientific explanation of behavior of linguistic incompetents or aspects of behavior which are not in fact subject to linguistic control as "mentalism," I was simply trying to capture a sense of that word which justifies what Dan Dennett has called the behaviorist's "gut intuition" (Dennett is talking about Skinner but all behaviorist psychologists, other than Tolman and his followers, share Skinner's reaction in this respect) "that the *traditional* way of talking about and explaining human behavior—in 'mentalistic' terms of a person's beliefs, desires, ideas, hopes, fears, feelings, emotions—is somehow utterly disqualified" (disqualified, that is, for use in a scientific theory). Psychological behaviorists, Skinner included, have never provided a rational justification for this prejudice against the use of common sense psychological language for scientific purposes. In particular they have never, to my knowledge, confronted the argument that if, as philosophical behaviorists

such as Wittgenstein and Ryle have claimed, most such language can be interpreted as talking about what the individual is disposed to publicly say and do, traditional objections to such language on the grounds that it is incurably subjective in its implications can no longer be sustained.

On the other hand, this "conceptual analysis" approach to what we are becoming accustomed to refer to as "folk psychology" reveals a number of other features of such language which arguably make it unsuitable for the purposes of scientific theory-construction. I have listed six such features in a recently completed article on this topic which is due to appear as another chapter in the same book (W. O'Donohue and R. Kitchener, Eds., *Psychology and Philosophy: Interdisciplinary Problems*) as my "Linguistic Behaviorism as a Philosophy of Empirical Science" [Place, 1996b] which I mentioned in my last letter:

- (1) the creation of bogus abstract entities by the process of "nominalizing" predicates and other nonsubstantival parts of speech,
- (2) the persistent use of adjectives with evaluative (good/bad) connotations,
- (3) the systematic evaluation of the content of other people's cognitive attitudes and judgments from the standpoint of the speaker,
- (4) the distortion of causal accounts of human action by the demand for a single scapegoat on whom to pin the blame when things go wrong,
- (5) the use of the metaphor of linguistic control when explaining behavior that is not subject to that type of control,
- (6) the unavoidable use of simile when describing private experience.

Of these the disqualification which is most important in relation to scientific explanations of language acquisition is 5, the use of language which presupposes linguistic competence on the part of the agent and linguistic control over the behavior to be explained. It goes without saying that unless you believe in something like Jerry Fodor's "language of thought," explanations of linguistic competence in these terms are viciously circular. It is this feature of

“mentalistic explanations” which, to my mind, constitutes the germ of truth underlying the “gut intuition” described by Dennett; hence (iii) in the list of behaviorist principles which I outlined in my last letter.

Chomsky (January 18, 1994)

What you say here has no resemblance, at least recognizable to me, to anything I said or hold. When I’ve discussed behaviorism, I’ve kept quite closely and explicitly to the doctrines espoused; to the texts themselves, in detail—explicitly, those who call themselves “behaviorists” (Quine, Skinner, Hull, etc.) and are so considered by others; I’ve never mentioned Watson (who you keep to) and I’ve also pointed out that these or very similar doctrines are often held by people who do not consider themselves behaviorists. You claim I’m creating a “convenient whipping boy”: I’d appreciate an example.

As for what you call my “contention,” I don’t even recognize it, let alone “contend” it. As for your “doctrine that an empirical science should be defined by . . .” it doesn’t matter how the dots are filled in. The particular empirical sciences are conveniences, nothing more; no one seeks to define them. No one tries to define “chemistry” as distinct from “physics.” True, MIT has a chemistry department and a physics department, for reasons of administrative efficiency and history. . . . You are addressing nonissues. I have no “contention” about them.

On your “third behaviorism principle, the issue of mentalism,” I’m afraid I’m now bewildered. You say what you do not believe, but not what you do believe about “the issue of mentalism,” apart from something about Dennett and “gut intuitions” of “behaviorists.” I’m afraid I can’t comment on that. As for Ryle’s critique, it’s largely beside the point, in my opinion, because (like the tradition generally) he seemed to believe that there is some

coherent notion of “physical” or “material.” Apart from that, he made serious errors of analysis, which I’ve discussed elsewhere, and had some useful particular insights, worth preserving (unfortunately, forgotten in much contemporary philosophy, notably at Oxford, sometimes in the most astonishing ways).

As for the use of common sense psychological language for scientific purposes, I agree with what you attribute to Skinner here. It’s not a “prejudice” to believe that common sense “physical” language (“the ball rolled down the hill and hit the ground”) is not very useful for scientific purposes. It would be a miracle if talk of beliefs, desires, etc., fared any better for these purposes than talk of the ground, rolling, etc. Wittgenstein and Ryle are not to the point here, even if we were to accept (I do not) their attitudes (not “arguments”; they don’t qualify as that) as to what language is “talking about.”

Place (May 16, 1994)

You challenge me to produce an example to illustrate my claim that you are selecting from the various components of the behaviorist position in such a way as “to create a convenient whipping boy.” You supply your own example by the names of those behaviorists you mention and those you ignore. You mention Quine, Hull, and Skinner. Of these Quine is not a psychologist and cannot possibly represent behaviorism as a movement in psychology, which is what it began as and has largely remained. His behaviorism, as far as I can see, is confined to his use of a bit of Skinner in his basic theory of language. You exclude Watson, the man who invented the term and defined its meaning, Lashley, his principal lieutenant in the early days, and Tolman who always described himself as a behaviorist, even if a “purposive” one. The one thing which these men had in common with Hull and Skinner is that they were all comparative psychologists who saw the experimental

study of animal behavior as a way of getting at the fundamental principles governing behavior, free from the complications of individual history, language, and human culture. The men you exclude are either, like Watson and Lashley, those who emphasize brain physiology as a way of explaining behavior or, like Tolman, someone who insisted on redefining our ordinary psychological terms in terms of objectively observable features of behavior, rather than abandoning them in favor of the language of "stimulus," "response," and "reinforcement."

I suggest that the effect of this selection is to create or, perhaps I should say, reinforce a stereotype of behaviorism in which three features predominate: (1) the repudiation of our ordinary mentalistic language for scientific purposes (the exclusion of Tolman), (2) the repudiation of explanations of behavior in terms of brain physiology (the inclusion of Skinner), (3) the analysis of behavior into mechanical stimulus-response connections (the inclusion of Hull). Of these, the first (although not explicit in Watson's writings) is certainly characteristic of most self-styled behaviorists with the notable exception of Tolman and Ryle. It is justified, as I have argued, by many features of our ordinary mental talk, but particularly by the use that is made in explanations of behavior of the device of quoting what the agent might be expected to say about the matter in hand. It now transpires that this is not a feature of the behaviorist position that you dispute. What you are not apparently willing to accept, however, is the conclusion that most behaviorists draw from this, namely that a radically different language is needed to describe and explain behavior for scientific purposes.

The second point (the repudiation of explanations of behavior in terms of brain physiology) is true, insofar as it is true of any behaviorist, only of Skinner. Even so, his position on the issue has been seriously misunderstood. What he is mainly concerned to assert

is his right to study the environment-behavior interface without prejudging the issue as to how the causal relation between the two is mediated by the brain. He does not deny either the importance of studying how the brain does it, though he thinks it's a job for the physiologist rather than the psychologist. I don't agree with him on this; but it's not a crucial point. The third point, the analysis of behavior into stimulus-response connections, was to all intents and purposes abandoned by Skinner in 1938 when he introduced the notion that the discriminative stimulus is the "occasion" under which a certain form of behavior (which he paradoxically still calls "a response") is "emitted" by the organism as a consequence of that kind of behavior having been "reinforced" under those conditions in the past.

Unfortunately traces of the older idea still survive in his later work, particularly in *Verbal Behavior*, where it underlies that appalling taxonomy of verbal operants classified according to the kind of stimuli that are alleged to "control" them.

You may say that you don't subscribe to the stereotype of behaviorism I have described. But you must concede that it is one that is widely held, particularly by those who have seen your 1959 review of Skinner's *Verbal Behavior* as its decisive refutation.

With regard to the issue of defining the scope of particular empirical sciences. I agree with you that such boundaries are often artificial and that preoccupation with them is a reflection of scientific immaturity. Nevertheless, it is clear that it has been and to some considerable extent remains a major issue within psychology and was the issue above all others that separated the behaviorists from introspective psychologists in the earlier part of this century. I would also claim that one of the things, perhaps the most important thing, that distinguishes a genuine empirical science from a pseudo-science is that the former takes as its subject matter, whether explicitly or implicitly,

a class of phenomena and leaves open the question as to how those phenomena are to be explained, while the latter, and I am thinking here of examples such as astrology and phrenology, are tied to a particular form of explanation. By that criterion, cognitive psychology insofar as it is tied to a particular way of explaining behavior is pseudo-science. So too is behaviorism, insofar as it conforms to the stereotype I have described or to the doctrines of a particular behaviorist such as Skinner. But that is not behaviorism as Watson originally conceived it, and as I would think it should be conceived. Looked at from this perspective, behaviorism is simply the claim that the behavior of living organisms is a legitimate field of study in its own right and should be approached as far as possible without any theoretical preconceptions about the right and wrong way of explaining that behavior. Other than that, it is for most behaviorists at least, the claim that you don't deny that as it stands, our ordinary mentalistic language is not a satisfactory medium for the construction of a genuine scientific theory in this area.

In conceding the latter point, you challenge me to say how I would handle the problem of replacing mentalism for scientific purposes. That is a long story and one on which my ideas are still in the process of development. Briefly stated, I distinguish within the psychological-neuroscientific domain three levels of description: (1) the neuro-synaptic level, (2) an intermediate level (the level of Edelman's "neuronal groups," the level of the network as a whole for the connectionist, the level of mediating behaviors for the behaviorists), and (3) the molar behavioral level.

At Level 1 in this scheme, I have no recommendations to make. The current language of neuroscience seems wholly adequate. At Level 2 I would use a combination of the language of associative learning theory and a selection of verbs drawn from our ordinary psychological language, those in particular

which *need not* take an embedded sentence in *oratio obliqua* as their grammatical object (verbs such as *attend to*, *imagine*, *expect*, *recognize*, *enjoy*, and a whole range of passive voice verbs of emotion). At Level 3 I would use a form of language based on Skinner's later writings, particularly his 1969 book *Contingencies of Reinforcement*. This would involve the concept of the three-term contingency, but characterized in terms of antecedent, behavior, and consequence, rather than, as Skinner characterizes it, in terms of stimulus, response, and reinforcement, and the distinction between rule-governed (or "rule-initiated" as I prefer to say) and contingency-shaped behavior. I would then use Skinner's concept of a rule as a "contingency-specifying (verbal) stimulus" to tie this way of talking into the ordinary notion of a means-end belief.

But since I have not yet managed to persuade *anyone* that this is the right way to proceed, I can hardly expect to persuade you.

I don't know what gives you the idea that Ryle seemed to believe that there is some coherent notion of "physical" or "mental" material. Having received my philosophical education at Oxford during the late 1940s when ordinary language philosophy was being expounded by Ryle, Austin, and my own philosophy tutor Paul Grice, I certainly came away from that experience, with the belief to which I have subscribed ever since, that it is the whole mental/physical distinction—both halves of it—which is incoherent, not just the mental half. Needless to say, I don't share your assessment of Ryle. Though I am critical of such things as his lack of sympathy for the scientific enterprise, his reluctance to be more systematic in developing his ontology, and his misunderstanding of causation, in other respects I am one of the few dyed-in-the-wool Ryleans still around. My view of dispositions is essentially Rylean, as is my conception of philosophical/linguistic methodology. But

here again, I can see that we are not likely to agree.

Suggested Readings

For further comments by Chomsky regarding behaviorism, see Chomsky (1959, pp. 26–58; 1965, p. 57, pp. 193–194, p. 204, p. 206; 1968/1972, pp. viii–ix, p. xi, p. 2, p. 4, pp. 25–26, pp. 35–38, p. 51, pp. 65–66, pp. 92–93, p. 118; 1973, pp. 104–150; 1987, pp. 158–182, pp. 436n–438n); 1988a, pp. 413–414; 1988b, pp. 137–138, pp. 161–162, pp. 165–166). For a discussion of linguistic behaviorism, see Place (1996a).

DISPOSITIONS

Chomsky (February 24, 1993)

[In your paper “Intentionality as the Mark of the Dispositional” (Place, 1996a)] you refer to “Chomsky’s principle that sentences are seldom repeated.” I can’t take credit for that. It goes back at least to Descartes, and is clear enough in Huarte in 1575 (which Descartes may have known). . . . [Additionally, in your paper] you say that it is a consequence of this principle that every sentence that a speaker constructs “is one amongst a range of possible manifestations of the disposition to construct . . . [an] utterance.” I don’t see how that’s a consequence of the principle; one may hold the principle (actually, more a trivial observation than a principle) and, consistently, deny that there is any such thing as “the disposition to construct . . . [an] utterance.” So it can’t be a consequence.

As I read Descartes, he would indeed have denied the existence of such dispositions. The Cartesians made a crucial distinction between what a person is “incited and inclined” to do (including, presumably, all dispositions) and what that person will indeed choose to do, irrespective of any dispositions, a distinction that crucially holds for language, indeed is the criterion for the existence of other minds.

Speaking for myself, I see no evidence for such dispositions, in any sense of the term “disposition” that has been given any sense, however vague. To select an example virtually at random, on reading . . . your paper [Place, 1991] on [the] analytic/synthetic [distinction], I was pleased to see that you traced the intension-extension distinction to Port Royal, rather than to Mill and Frege as is commonly done; and I can easily imagine how one could disagree with Quine that there is an intimate connection between the analyticity of a proposition and the intensions of its terms, for example, by questioning whether the terms of a proposition have fixed and absolute intensions, that is, whether this is the proper picture of language at all (I doubt that it is). But I can’t see any meaningful sense in which I had a disposition to construct the preceding sentence on reading that page. That seems to me just hand-waving, borrowing the term “disposition” that has a relatively clear sense in other contexts, and applying it here, metaphorically, though it has no sense in this context. It’s the kind of thing Quine does all the time, for example; but improperly.

Anyway, I can’t take credit for the “principle” (which is more an observation than a principle in any event), and I don’t think the consequence follows.

Place (June 24, 1993)

With regard to what you say in your letter, I think you are being much too modest in refusing to accept credit for the observation that sentences are seldom repeated and are typically constructed anew on each occasion of utterance. It may well be that others have drawn attention to the phenomenon; but I don’t think that you can reasonably deny that it was you who made the *explanation* of that phenomenon the central issue for contemporary linguistic theory.

I am puzzled by your denial that your “observation” that sentences are

constructed anew on each occasion of utterance, when combined with the observation that there are indefinitely many sentences which will express the same "thought" in Frege's (1918/1956) sense of that word, implies that the utterance of a sentence is preceded as a matter of psychology by a disposition to construct a range of possible sentences of which the actual sentence constructed and uttered is only one. That the constructor of a sentence is confronted with a number of choices between different ways of "putting" what she wants to say is, I should have thought, a matter of common introspective observation. As Wundt (1907) puts it, it is a common experience that

a thought is not first formed while one speaks a sentence, . . . it already stands as a whole in our consciousness before we begin to fit words to it. With this whole there is, nevertheless, present at the focus of consciousness none of the verbal or other representations which form during the development and the linguistic expression of the thought; but only at the moment when we develop the thoughts are their separate parts successively lifted to clear consciousness. (Wundt, 1907, p. 349, quoted and translated by Humphrey, 1951, pp. 110–111)

Clearly the range of possible sentences expressing the same thought which the individual speaker has, as it were "on the tip of her tongue" and amongst which she makes her choice is narrower or broader depending upon her degree of linguistic competence, but is in any case vastly more restricted than the complete range of possible sentences in any natural language with the conceptual resources needed to express it which comprise the thought in Frege's sense.

It seems to me that to deny, as you seem to do, that such a disposition precedes the construction of a sentence or, if it does, that it stands in any kind of causal relation to the sentence that is actually constructed is to deny any sense to the common expression whereby a speaker is said to "choose her words carefully." To say, as you seem to want to do, that it makes no sense to speak of sentences the speaker

might have chosen to construct on that occasion, but did not, smacks to my mind of the worst excesses of positivism, and Quinean extensionalism.

You may be right that, strictly speaking, to hold (a) that sentences are generally constructed anew on each occasion of utterance, and (b) that even within a given natural language for any thought there are indefinitely many sentences which could be used to express it, is not inconsistent with denying that a disposition to construct a range of possible sentences all of which express the same thought as the sentence ultimately constructed typically precedes and guides the process of sentence construction; but that the three doctrines are natural bedfellows which mutually support and make sense of one another cannot, I think, be denied.

That there is some sort of misunderstanding here on your part is suggested when you say "I can't see any meaningful sense in which I had a disposition to construct the preceding sentence on reading that page." As I point out in arguing against Martin and Pfeifer's attempt to generate a referentially opaque description of a physical disposition, open-endedness (Anscombe's "indeterminacy") is of the essence of a disposition. You did not have a disposition to construct *that* sentence. What you had was a disposition to construct a range of possible functionally equivalent sentences of which *that* sentence was only one.

Chomsky (July 22, 1993)

If we deprive the notion "disposition" of any substantive meaning, as Quine and many others do, we can then safely conclude that a disposition precedes the construction of a sentence. If we use the term with anything like its actual meaning, the claim that dispositions precede sentence construction is about as convincing as the claim that objects fall to their natural place. True, the claim conforms to some (it seems largely empty) doctrines about

causation of behavior, but it has no other merits, to my knowledge, and no supporting evidence. Also, for what this matters (not much), it does not conform to intuition, at least mine. Thus, I have no intuition that my production of this sentence was preceded by a probability distribution over possible utterances relative to the given circumstances (except, of course, for the distribution that assigns to this utterance, and a vast number of others, a probability that is vanishingly small). As for people choosing what is on the tip of the tongue, or choosing their words carefully, surely such phenomena exist, but they leave the general claim unsubstantiated; at most, they indicate that in these particular and in fact rather marginal cases, the speaker may have been disposed to say such and such—if that turns out to be the right way to approach an explanatory account.

I don't say it "makes no sense to speak of sentences the speaker might have chosen." Surely it makes perfect sense. Thus, on the current occasion, I might have chosen any of an infinity of expressions, an infinite subset of them being perfectly well-formed sentences of my language (others, perhaps the ones I end up producing, not being so, for various reasons). The problem is not that the notion makes no sense, leading to your Quinean rebuke, but that there is no known way of giving it any *interesting* content. It's nothing more than a slogan, introduced to make the discussion look somehow scientific. Slogans have sense; they are just not very helpful.

There are real problems in accounting for behavior, and we are only led on false tracks, and given a false sense of security and accomplishment, by declaring without credible argument that behavior is "chosen" or "caused" or "preceded by dispositions" in any sense of these terms that is at all understood.

Place (August 21, 1993)

I agree that there has been a lot of loose and empty talk about dispositions

and their role in causation. But I trust that I can be exonerated from the accusation of having contributed to it. The charge to which I will plead guilty is that of not having published my account of the matter in publications which are readily accessible to the philosophical community in the English-speaking world. My paper "Causal Laws, Dispositional Properties and Causal Explanations" [Place, 1987] was published in *Synthesis Philosophica* (the international edition of the Serbo-Croat journal *Filozofska Istrazivanja*) in 1987. I sent a copy to David Armstrong which provoked a correspondence which led eventually to our "Debate on Dispositions and Their Role in Causation," which appeared in the Austrian philosophy journal *Conceptus* (Armstrong & Place, 1991). We are hoping that, with the addition of a contribution from C. B. Martin, this will eventually be published in book form by Routledge and Kegan Paul [see Armstrong, Martin, Place, & Crane, 1996]. Copies of both these publications are enclosed.

Chomsky (November 8, 1993)

Here there are many interesting issues, on many of which I am sympathetic to your views (e.g., 20 years ago I argued that Kripke was misunderstanding his intuitions, which, insofar as they are valid, have to do with *de dicto* necessity (in the relevant cases)). But I don't feel that these issues, however resolved, have anything much to say about language and its use, because the underlying conception of these matters is (in my view) misguided. I won't repeat.

Place (December 29, 1993)

It seems that there is much agreement here; but also disagreement in that you see no relevance of these issues in relation to "language and its use." I see the concept of a dispositional property as an integral part of the concept of causation, indeed as the "cement," to use Hume's metaphor,

which binds cause to effect. As such, dispositional properties are omnipresent and omnipotent. Without them, nothing changes; nothing persists. Language is no exception in this respect. After all what is linguistic competence if not a dispositional property of the language user.

Chomsky (January 18, 1994)

The problem is that no one, to my knowledge, has offered any notion of "disposition" that has any bearing on the issues. Thus, if a disposition involves a probability distribution over possible actions in particular circumstances, then the only dispositions in language use are to say "hello!" when you walk into a room, and a few other such things. Linguistic competence is most definitely *not* a dispositional property of the language user, in any sense of "dispositional property" that has ever been presented. That's akin to the utterly failed effort to reduce knowledge to some kind of ability. These are just mantras, I'm afraid, used to make philosophers feel comfortable, signifying nothing.

Place (May 16, 1994)

Having spent much of the past four years writing a book on this topic jointly with my old sparring partners, David Armstrong and Charlie Martin, I can't agree that no one has given an adequate account of dispositions. Though, since our book has not yet appeared and since we all take up different positions on the issue, I can't claim that a final agreed solution to the problem has been found, let alone announced to the world. However, what is quite clear is that, despite our disagreements on other aspects of the matter, none of us would agree with you in thinking (a) that the only linguistic dispositions there are, are dispositions to utter the same word or string of words on certain types of occasion, (b) that linguistic competence is not a disposition, (c) that to know something is not an ability and hence not a disposition, and (d)

that to say of something that it is a disposition is mere hand waving. The three of us would all agree that whatever else it is, a disposition is that whose existence *qua* property of its owner makes true a subjunctive conditional of the form "If at any time such and such conditions were to be fulfilled, an event or state of affairs of a certain generic type would or would very probably occur or exist?" Any such event or state of affairs if it were to occur or exist would constitute a *manifestation* of the disposition in question. We would also all agree that the range of possible manifestations which characterize a particular disposition is always open-ended. Thus to take your case of the disposition to say "Hello!" there are clearly indefinitely many different tones of voice in which that combination of phonemes can be uttered, each with its different nuance in terms of warmth, irony, etc. We would also all agree that the breadth and variety of possible ways in which a disposition can manifest itself will vary from disposition to disposition. In the case of a disposition such as linguistic competence, which is a combination both of the speaker's ability to construct and produce intelligible sentences and the listener's ability to construe them, this breadth and variety of possible manifestations is enormous. In other cases such as the brittleness of glass it is restricted to the range of possible ways in which the object could crack, break, or shatter. In the case of knowing, the nature of the disposition varies depending on the kind of grammatical object taken in the verb "know." In the typical case where the grammatical object is an embedded interrogative sentence introduced by an interrogative pronoun, to know is to be able to provide a correct answer to the question specified by the embedded interrogative and act accordingly. How wide is the range of possible correct answers will vary from question to question.

This open-ended subjunctive modal character of dispositions makes them,

of course, extremely offensive to those brought up, as I gather you are, in the Quinean tradition, which takes an extensional logic such as the predicate calculus as the measure of all things logical, philosophical, and linguistic. Since I believe that dispositions are everywhere in nature and are an essential ingredient in every causal relation, I take this to show the total inadequacy of an extensional logic as the basis for any empirical science. It also leads me to think that the attempt to formalize linguistics is radically misconceived; but that is an opinion which I cannot expect you to share.

Suggested Readings

For a further elaboration of Place's views on dispositions, see Place (1987, 1996a). Also see Armstrong and Place (1991) and Armstrong et al. (1996). For an important historical antecedent view, see Ryle (1949).

REFERENCES

- Armstrong, D. M., Martin, C. B., Place, U. T., & Crane, T. (Eds.). (1996). *Dispositions: A debate*. London: Routledge.
- Armstrong, D. M., & Place, U. T. (1991). A debate on dispositions: Their nature and their role in causation. *Conceptus, Band 25, No. 66*, 3–44.
- Block, N. (1990). The computer model of the mind. In D. N. Osherson & E. E. Smith (Eds.), *Thinking: An invitation to cognitive science* (Vol. 3, pp. 247–289). Cambridge, MA: MIT Press.
- Broadbent, D. E. (1991, December). *Planning and opportunism*. Address to the London Conference of the British Psychological Society, City University, London.
- Chomsky, N. (1959). Review of B. F. Skinner's *Verbal Behavior*. *Language*, 35, 26–58.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1972). *Language and mind* (enlarged ed.). New York: Harcourt Brace Jovanovich. (Original work published 1968)
- Chomsky, N. (1973). *For reasons of state*. London: Fontana.
- Chomsky, N. (1980). *Rules and representations*. New York: Columbia University Press.
- Chomsky, N. (1986). *Knowledge of language: Its nature, origin and use*. New York: Praeger.
- Chomsky, N. (1987). *The Chomsky reader*. New York: Pantheon.
- Chomsky, N. (1988a). *Language and politics*. Cheektowaga, NY: Black Rose Books.
- Chomsky, N. (1988b). *Language and problems of knowledge: The Managua lectures*. Cambridge, MA: MIT Press.
- Chomsky, N. (1993). *Language and thought*. Wakefield, RI: Moyer Bell.
- Chomsky, N. (1997). *Language and mind: Current thoughts on ancient problems* (part 1). Available on line at kcen.ru/tatlen/science/fccl/papers/chomsky1.htm.
- Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge, UK: Cambridge University Press.
- Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. New York: Harper and Row.
- Churchland, P. M. (1988). *Matter and consciousness* (rev. ed.). Cambridge, MA: MIT Press.
- Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group selection*. New York: Basic Books.
- Frege, G. (1956). The thought: A logical enquiry (A. Quinton and M. Quinton, Trans.). *Mind*, 65, 289–311. (Original work published 1918)
- Gehrz, R. D., Black, D. C., & Solomon, P. M. (1984). The formation of stellar systems from interstellar molecular clouds. *Science*, 224, 823–830.
- Helm, G. (1993, June). *Computers can think! A strange proof and its implications for cognitive science*. Paper presented to the Conference on Connectionism and the Philosophy of Mind, Bled, Slovenia.
- Herrnstein, R. J., Loveland, D. H., & Cable, C. (1976). Natural concepts in pigeons. *Journal of Experimental Psychology: Animal Behaviour Processes*, 2, 285–302.
- Humphrey, G. (1951). *Thinking: An introduction to its experimental psychology*. London: Methuen.
- Jordan, M. I. (1986). Attractor dynamics and parallelism in a connectionist sequential machine. *Proceedings of the eighth annual meeting of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Lashley, K. S. (1930). Basic neural mechanisms in behavior. *Psychological Review*, 37, 1–24.
- Lashley, K. S. (1938). The mechanism of vision: XV. Preliminary studies of the rat's capacity for detail vision. *Journal of General Psychology*, 18, 123–193.
- Lewontin, R. C. (1990). The evolution of cognition. In D. N. Osherson & E. E. Smith (Eds.), *Thinking: An invitation to cognitive*

- science* (Vol. 3, pp. 229–246). Cambridge, MA: MIT Press.
- Lycan, W. G. (1990). *Mind and cognition: A reader*. Oxford: Blackwell.
- Matthews, P. H. (1993). *Grammatical theory in the United States from Bloomfield to Chomsky*. Cambridge, UK: Cambridge University Press.
- Moerk, E. L. (1983). *The mother of Eve as a first language teacher*. Norwood, NJ: Ablex.
- Pearce, J. M. (1988). Stimulus generalization and the acquisition of categories by pigeons. In L. Weiskrantz (Ed.), *Thought without language* (pp. 132–155). Oxford: Clarendon Press.
- Pearce, J. M. (1989). The acquisition of an artificial category by pigeons. *Quarterly Journal of Experimental Psychology*, 41B, 381–406.
- Pinker, S. (1994). *The language instinct*. New York: William Morrow.
- Pinker, S., & Bloom, P. (1990). Natural language and natural selection. *Behavioral and Brain Sciences*, 13, 707–784.
- Place, U. T. (1956). Is consciousness a brain process? *British Journal of Psychology*, 47, 44–50.
- Place, U. T. (1987). Causal laws, dispositional properties and causal explanations. *Synthesis Philosophica*, 2(3), 149–160.
- Place, U. T. (1991). On the social relativity of truth and the analytic/synthetic distinction. *Human Studies*, 14, 265–285.
- Place, U. T. (1992). Eliminative connectionism: Its implications for a return to an empiricist/behaviorist linguistics. *Behavior and Philosophy*, 20, 21–35.
- Place, U. T. (1996a). Intentionality as the mark of the dispositional. *Dialectica*, 50, 91–120.
- Place, U. T. (1996b). Linguistic behaviorism as a philosophy of empirical science. In W. O'Donohue & R. F. Kitchener (Eds.), *The philosophy of psychology* (pp. 126–140). London: Sage.
- Reeke, G. N., & Edelman, G. M. (1988). Real brains and artificial intelligence. *Dædalus*, 117, 143–173.
- Ryle, G. (1949). *The concept of mind*. London: Hutchinson.
- Schoneberger, T. (1996, May). *Chomsky in retreat*. Paper presented at the 22nd annual convention of the Association for Behavior Analysis, San Francisco.
- Skinner, B. F. (1957). *Verbal behavior*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1969). *Contingencies of reinforcement: A theoretical analysis*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1971). *Beyond freedom and dignity*. New York: Bantam/Vintage.
- Wundt, W. (1907). Über ausfragexperimente über die methoden zur psychologie des denkens. *Psychologische*, 3, 301–360.